Utility Maximizing Node Activation Policies in Networks of Partially Rechargeable Sensors

Neeraj Jaggi, Ananth Krishnamurthy and Koushik Kar

***Abstract*—We address the problem of optimal node activation in a sensor network, where the optimization objective is represented by a time-average generalized coverage metric. We consider a scenario where each sensor is rechargeable, can hold up to K quanta of energy, and can be activated even if it is partially recharged. If the recharge and/or discharge process is random, the problem of optimal sensor activation in this context is a complex stochastic decision question. For the case of identical sensor coverages, we show that there exists a simple *threshold policy* that is *asymptotically optimal* with respect to K, i.e., the performance of the this threshold policy approaches the optimal performance as K becomes large. We then extend this approach to a general sensor network, where coverage areas of different sensors could have full, partial or no overlap with each other. We demonstrate through simulations that a local information based threshold policy, with an appropriately chosen threshold, achieves a performance that is very close to the global optimum. We also show, through analysis and simulations, that spatial correlation in the discharge process worsens system performance.**

I. Introduction

Significant advancements in hardware technologies, and growing application requirements, have fueled a considerable interest in ad-hoc and sensor networks in recent years. It is envisioned that in the next few decades, these networks will be used for communication in infrastructure-less regions, military and relief operations, intra-community communication, environmental and health monitoring, and a vast number of other applications. A large number of such networks consist of battery-powered nodes, which implies that these networks are also heavily constrained in terms of energy [13]. Typically, a sensing device can remain powered on (and be sensing) only for a limited amount of time, until it runs out of battery energy [13]. In many scenarios, however, sensors may be equipped with rechargeable batteries, but recharging is often a very slow process, possibly influenced by random environmental factors like the intensity of sunlight etc. Typically, the average rate of recharging will be significantly less than the average energy discharge rate during the sensing period. As a result, a sensor could need to spend most of its lifetime in the "off" state, when it is not sensing, but only recharging. These factors motivate redundant deployment of sensors to cover the area of interest, so that more number of these sensors would remain charged (and hence can be used for sensing) at any given time. In the following, we consider such a scenario, where sensors have been deployed redundantly in the area of interest.

N. Jaggi and K. Kar are with the Department of Electrical, Computer and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180, USA (email: jaggin@rpi.edu, koushik@ecse.rpi.edu). A. Krishnamurthy is with the Department of Decision Sciences and Engineering Systems, Rensselaer Polytechnic Institute, Troy, NY 12180, USA (email: krisha@rpi.edu).

There is a large body of work on energy-efficient MAC and adaptive wakeup of sensors, although these works only consider energy-constrained, but non-rechargeable sensors. A discussion on the importance of energy management in adhoc and sensors networks, along with a description of various performance objectives, is outlined in [13]. Energy-efficient MAC protocols are studied in [5], [6], [15]. The problem of minimizing power consumption during idle times is addressed in [4], [9], [14]. In [3], the authors use occupancy theory to analyze the effect of switching off idle nodes on the network lifetime. In [11], the effects of power conservation, coverage and cooperation on data dissemination is investigated for a particular data sharing architecture.

Unlike previous work, in this paper we consider sensors that are rechargeable. We assume that sensors are recharged continuously, according to a random process. On the other hand, sensors *activate* (i.e., participate in sensing and transmission) themselves periodically, and are discharged during the activation period, again according to some random process. The decision question that we address is when and which sensors should be activated (or "switched on") so that the time-average system utility is maximized.

The paper is structured as follows. In the next section, we formulate the sensor activation problem as a stochastic decision question. We outline our basic approach and contribution in Section III. In Section IV, we consider a system of sensors with identical coverages, and show that simple threshold policies can achieve asymptotically optimal performance in this scenario. In Section V, we extend the threshold activation policy outlined in Section IV to a more general scenario where sensor coverages may not overlap completely, and evaluate its performance through simulations.

II. Formulation

We assume that battery recharge occurs in units of an energy *quantum*. Each sensor is modeled as a K-quantum bucket, i.e., a sensor battery can hold at most K quanta. Quanta arrive at the sensor according to a random process. During its activation period of a sensor, these accumulated quanta are used up one by one. Thus, a sensor can be modeled as a finite-buffer quantum-queue, as shown in Figure 1, with a buffer capacity of K. Note that this queue is discharged of quanta only when the sensor is activated. Therefore, the quantum service (discharge) process depends on the chosen activation decision policy. More details on the quantum arrival (recharge) and service (discharge) process models are provided in later in this paper.

We assume that the performance of the system is characterized by a continuous, non-decreasing, strictly concave function

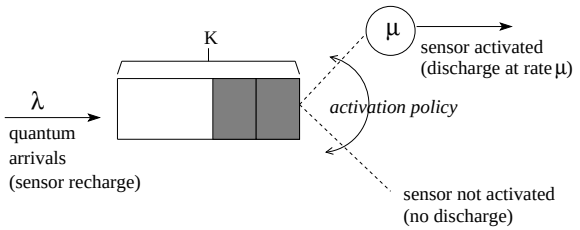

Fig. 1. Quantum-queue model of a sensor

U satisfying $U(0) = 0$. More specifically, $U(n)$ represents the utility derived *per unit area, per unit time*, from n active sensors covering an area. Note that different sensors can be located at different points in the overall physical space of interest, and the coverage patterns of different nodes can be different. Therefore, the coverage areas of different sensors will typically be different. This implies that at any time, utilities in different parts of the area of interest can differ significantly from one another.

Note that the strict concavity assumption merely states the fact that the system has diminishing returns with respect to the number of active sensors. As an example of a practical utility function, consider the scenario where each sensor can detect an event with probability p_d. If the utility is defined as the probability that the sensing system is able to detect an event, then $U(n) = 1 - (1 - p_d)^n$, where n is the number of sensors that are active. Note that this utility function is strictly concave, and satisfies $U(0) = 0$.

We are interested in maximizing the time-average utility of the system. Let $\mathcal{A}$ denote a generic area element in the physical space of interest. Let $n_P(A, t)$ denote the number of active sensors that cover area element A at time t, under activation policy P. The time-average utility under policy P, is given by

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t \int_{\mathcal{A}} U(n_P(A,t))\, dA\, dt\,. \tag{1}$$

In Euclidean coordinates system, $dA = dx\ dy$, and $n_P(A, t) = n_P(x, y, t)$, in the above expression. The decision problem that we consider in this paper is that of finding P so that the objective function in (1) is maximized.

Clearly, the sensors with no energy (i.e., sensors whose quantum-queue is empty) cannot be activated. Therefore, our decision problem is that of determining how many, and which, sensors to activate at any time, from the set of *available* sensors (i.e., sensors with non-zero energy). Note that if we activate more sensors, we gain utility in the short time-scale. However, if the number of active sensors is already large, since the utility function exhibits diminishing returns, we may want to keep some of the available sensors "in store" for future use. In fact, the performance results shown later in this paper justify this intuition.

III. Methodology and Contribution

The random nature of recharging and discharging process implies that the activation question outlined above is a stochastic decision problem. Although under specific cases the optimal policies may be formulated as semi-Markov decision problem, determining optimal policies can be computationally prohibitive. Since sensors are energy constrained we seek policies that can be implemented in a distributed manner with minimal information and computational overhead. Therefore, in this paper, we focus on simple *threshold* policies (defined precisely later in the paper) and examine their performance. Interestingly, as we show later in this paper, a threshold policy leads to near-optimal performance, if the threshold is chosen appropriately.

To simplify the analysis and obtain fundamental performance insights, we examine the performance of threshold policies for a system of sensors whose coverage areas are identical. In this case the objective function in equation (1) reduces to a single integral over the time domain. Under Markovian assumptions we derive tight bounds on the performance of threshold policies for two different lifetime correlation models of the sensor nodes. These performance bounds motivate our study of the performance of threshold based policies in a very general network setting, where coverage areas of different sensors could have partial, complete or no overlap with each other. Results from our extensive simulation studies show that even in this case, the performance of threshold activation policies is still very close to the best achievable performance.

The optimal node activation problem for rechargeable sensor networks was first addressed by Kar *et al.* in a recent work [8]. In [8], for the case of identical sensor coverages, the authors have formulated the node activation problem as a complex semi-Markov decision problem (MDP), and provided a solution approach to this problem. This solution approach being computationally hard, the authors have studied the performance of threshold policies, and shown that the best threshold policy achieves a performance within $\frac{3}{4}$ of the best achievable performance. Threshold policies are also shown to work well in a general network scenario with non-identical sensor coverages.

In this paper too, we study the performance of threshold policies for the optimal node activation problems, but our results differ from, and enhances, the results in [8] in several aspects. Firstly, in [8], the authors consider a exponential model for the recharge and discharge times; our finite-buffer quantum queue model is a much more realistic model of the recharge/discharge process. Secondly, in order to simplify the analysis, the system model in [8] makes the assumption that sensors can only be activated only when they are fully recharged. However, there is no reason why sensors with partially recharged batteries cannot take part in the sensing process. In this paper, we allow partially recharged sensors to be activated, and show that this could improve the system performance significantly over what was observed in [8]. The consideration of partially recharged node activation is a key aspect that distinguishes our work from [8]. Finally, the performance bounds that we obtain are significantly stronger than those observed in [8]. We show that threshold policies can achieve a performance that is *asymptotically optimal*, with respect to the buffer capacity K. Since K is expected be fairly large (K represents the granularity on the energy recharge/discharge process) in most practical scenarios, the difference between performance of the threshold policy and that of the optimal policy is insignificant in almost all cases. On the other hand, the performance bound obtained in [8] for threshold policies is only $\frac{3}{4}$ of the optimum. Note that this performance improvement, and the near-optimality of threshold policies, are only possible because of the activation of partially

recharged sensors.

It is worth noting here that the threshold policies that we study can be implemented in a large-scale sensor network in a distributed manner, with only local information. Therefore, our results show how performance close to the global optimum (which is computationally difficult to obtain) can be achieved with a simple activation policy using only local information.

IV. Activation Policies with Identical Sensor Coverages

In this section, we address our node activation decision problem for a system of sensors whose coverage areas overlap completely with each other.

A. Problem Definition

We consider a system of N sensors covering the same area. We model the quanta arrival (recharge) process at each sensor by a poisson process with rate λ. Note that in a realistic sensing environment, the discharge time of a quantum could depend on various random factors. For instance, sensors can transmit information (resulting in energy usage) on the occurrence of "interesting" events, which may be generated according to a random process. Therefore in our system model, we assume that the quantum discharge times are random. More specifically, we assume that the discharge time of each quantum is exponentially distributed with mean $\frac{1}{\mu}$. In the simulations, however, we also study the performance with deterministic discharge and recharge times. Note here that we assume $\mu \geq \lambda$, which captures the fact that the discharge rate of a sensor in the active state is no less than the recharge rate.

We assume that there is no energy discharge when a node is not active. However, note that although a sensor can power itself off when it is not active, it has to wake up periodically and exchange messages with its neighbors to keep track of the system state in its neighborhood (so as to decide whether to activate itself or not). Therefore, in reality, we would expect that energy will be drained even when a sensor is not active (as long as it has non-zero energy), but probably at a fairly steady rate. However, since the energy discharge rate in the non-active state can be expected to be much slower than the discharge rate in the active state, it is not considered in our analysis.

Let $n_P(t)$ denote the number of sensors in the active state at time t under policy P. Since the coverage areas of all sensors are completely overlapping, the optimization problem can be posed as that of finding a policy P that minimizes $\bar{U}(P)$, where $\bar{U}(P)$ is defined as

$$\bar{U}(P) \quad = \quad \lim_{t\to\infty} \frac{1}{t} \int_0^t U(n_P(t))dt \; . \tag{2}$$

We assume that activation decisions can be taken at any instant of time. As we argue later in the paper, these decisions need to be taken only when some active sensor runs out of energy, or when some sensor with zero energy becomes available by gaining a quantum through recharging. It is worth noting here that although we will address our decision question in the identical coverage case from the perspective of a centralized decision maker, the decision policy that we develop can easily be implemented in a decentralized manner.

B. Discharge Process Models

As we discuss later in more detail, the performance of decision policies depend considerably on how the discharge processes of sensors are correlated. We consider two different correlation models of the discharge/recharge times of the different sensors:

- *Independent Discharge (ID) process model:* In this model, the quantum discharge times of the active sensors are independent of each other.
- *Correlated Discharge (CD) process model:* In this model, all active sensors get discharged of a quantum at the same time.

The two correlation models can be practically motivated in the following way. To motivate the ID model, consider a scenario where data transmission (on the detection of interesting events) is the primary mode of energy expenditure, as is often the case in practice. Moreover, assume that the detection of an event and/or the subsequent data transmission is a random variable. This could happen, for instance, if a certain failure probability is associated with the detection of an event at each sensor, or if an event is reported only by one or a subset of the sensors (randomly chosen) detecting the event (to avoid redundant event reporting). In such scenarios, the system is better represented by the ID model. To motivate the CD model, consider a scenario where all active sensor report data on a regular basis, or the data reporting is so infrequent that most of the energy expenditure occurs in sensing and processing. In these cases, the active sensors will all get discharged at the same rate, and therefore, the system is better modeled by the CD model.

Note that these two models represent two extreme forms of correlation, and real-life situations can be expected to fall in between these two extremes. We will argue that our solutions perform well with respect to both of these extreme forms of correlation. Therefore, our solutions are expected to perform well in intermediate discharge correlation scenarios as well.

Note that the recharge process can exhibit a certain degree of spatial correlation as well. Since we discuss the case of identical sensor coverages in this section, we will assume that the recharge processes at the sensors are perfectly correlated, i.e., quanta arrive at all the sensors simultaneously. Since the sensors have identical coverages, they can be assumed to be located close to one another; thus, the high correlation between the recharge processes at different sensors are justified. Our results can also be extended to the case of independent recharge processes, and certain intermediate scenarios. The fundamental observations in those cases are similar to those in the perfectly correlated recharge process case addressed in this paper, and those are therefore omitted for the sake of brevity.

C. Upper Bound on the Optimal Time-average Utility

Since the optimal time-average utility is difficult to compute, we obtain an upper bound on it. In the next subsection, we will evaluate the performance of our proposed policies with respect to this bound. The proof of the following result is provided in the Appendix.

Lemma 1: The time-average utility under *any* node activation policy is upper-bounded by $U(\frac{N}{\gamma})$.

Note that this results holds for both ID and CD process models.

D. Performance Analysis of Threshold Activation Policies

Next we define the class of threshold activation policies, and analyze their performance.

1) Threshold Activation Policies: Note that the set of all possible activation policies can be very large, and the structure of these policies can be very complex. Therefore, determining the optimal activation policy for the ID and CD models is computationally very intensive. Therefore, we focus primarily on threshold activation policies. A threshold activation policy with parameter m, is characterized as follows: an available sensor (i.e., a sensor with non-zero energy) s is activated if the number of active sensors does not exceed m after s is activated; otherwise, s is not activated. In other words, a threshold policy with parameter m tries to maintain the number of active sensors as close to m as possible. Note that with such a policy, the number of active sensors can never exceed m.

2) Scheduling Models: The performance of the threshold policies, and their analyses, depend on how the active sensors are chosen from the set of available sensors. Recall that for a threshold with parameter m, at most m sensors are activated at any given time. Most common scheduling disciplines (work-conserving as well as non work-conserving) can be used in selecting the sensors to be activated. In the following, we analyze the performance of threshold policies for a *Longest Undischarged-time First (LUF)* scheduling policy. In LUF, at any decision instant, the sensor that is activated first is the one which has not been discharged for the longest time, amongst the set of available sensors. In other words, at each decision instant, available sensors are ordered according to the last time they were active, and the first sensor in the increasing order is chosen for activation. Note that the LUF policy tries to achieve fairness amongst sensors while choosing them for activation.

3) Performance Bounds: Now we derive worst-case bounds on the performance of threshold policies (with LUF scheduling) with respect to the optimal policy, for the two correlation models. We define $\gamma = \frac{\mu}{\lambda} \geq 1$. For simplicity of exposition, we assume γ is an integer, and N is divisible by γ (i.e., $\frac{N}{\gamma}$ is an integer), although our results can be generalized to cases where these assumptions do not hold.

First, let us consider the CD model. Let $V_C(m)$ denote the time-average utility for the CD model, with LUF scheduling, at threshold m.

Theorem 2: Let $m^* = \frac{N}{\gamma}$. Then $V_C(m^*)$ is lower-bounded by $\frac{K}{K+1}U(\frac{N}{\gamma})$, i.e.,

$$V_C(m^*) \geq \frac{K}{K+1}U(\frac{N}{\gamma}) \, .$$

The proof outline of Theorem 2 is provided in the Appendix. Theorem 2, together with Theorem 1, implies threshold policy with a threshold of $m^* = \frac{N}{\gamma}$ is within a factor of $\frac{K}{K+1}$ of the optimal performance over all policies.

For the ID model too, we can show a result similar to the one stated above. Let $V_I(m)$ denote the time-average utility for the ID model, with LUF scheduling, at threshold m.

Theorem 3: Let $m^* = \frac{N}{\gamma}$. Then $V_I(m^*)$ is lower-bounded by $\frac{K}{K+1}U(\frac{N}{\gamma})$, i.e.,

$$V_I(m^*) \geq \frac{K}{K+1}U(\frac{N}{\gamma}) \, .$$

The proof of the above result is in the Appendix. Theorem 3, together with Theorem 1, implies threshold policy with a threshold of $m^* = \frac{N}{\gamma}$ is within a factor of $\frac{K}{K+1}$ of the optimal performance over all policies.

Note that the performance ratio $\frac{K}{K+1}$ approaches unity as K becomes large. Therefore the threshold activation policy with a threshold of $m^* = \frac{N}{\gamma}$ is asymptotically optimal with respect to K, for both the ID and CD models.

Next we compare the performance of threshold policies for the ID and CD models. Consider a threshold activation policy with a threshold of m, and assume that N is a multiple of m. First, for the ID model, we define a *group Longest Undischarged-time First (g-LUF)* policy as follows. Divide the N sensors into m groups, each group consisting of $\frac{N}{m}$ sensors. Under the g-LUF policy, at any decision instant, at most one sensor is activated from each group. Within a group, the sensor to be activated is chosen (from the set of $\frac{N}{m}$ sensors of that group) based on a LUF scheduling policy. Note that if none of the sensors in any group is available (i.e., all the sensors in that group have run out of energy), then the threshold of m cannot be met, even if the all the sensors in the other groups are available. Thus, g-LUF is a non work-conserving scheduling discipline. Let $V_I^g(m)$ denote the time-average utility for the CD model, with g-LUF scheduling, at threshold m.

Let S_N denote the set of all positive factors of N, i.e., all positive integers which divide N. Then the following result (proof in Appendix) shows that the threshold policy under ID model and g-LUF scheduling performs no worse than a threshold policy under CD model and LUF scheduling.

Theorem 4: For any $m \in S_N$, $V_I^g(m)$ is no less than $V_C(m)$, i.e.,

$$V_I^g(m) \geq V_C(m) \quad \forall m \in S_N.$$

Note that under the ID model, the performance with LUF scheduling (which is work-conserving) is intuitively expected to be better than that with g-LUF scheduling. Therefore, with LUF scheduling, the performance under ID model is expected to be better than that under the CD model. Proof of this fact remains an open question; however, we support this observation through simulation results in the next subsection.

E. Simulation Validation

In this section, we report results from numerical experiments on the performance of threshold policies for the ID and CD models under different parameter settings. The utility function used is $U(n) = 1 - (1 - p_d)^n$, with $p_d = 0.1$. We report the results for $N = 16$, $K = 10, 100$, and $\gamma = 2, 4$; the results obtained for other values of these parameters are similar in nature. For each parameter setting, we compare the time-average utility of the system for different values of the threshold (m). The scheduling policy used in the simulations for both the ID and CD models is LUF.

Figures 2 - 5 show the time-average utilities $V_I(m)$ and $V_C(m)$ along with $U(\frac{N}{\gamma})$, the upper bound on the maximum achievable time-average utility. The figures also show $\frac{K}{K+1}U(\frac{N}{\gamma})$, the lower bound on the achieved utility for threshold policy at a threshold of $m^* = \frac{N}{\gamma}$. The figures show that the performance of the threshold policy is maximized at the threshold of $m^* = \frac{N}{\gamma}$ or close to it. Moreover, the time average utility

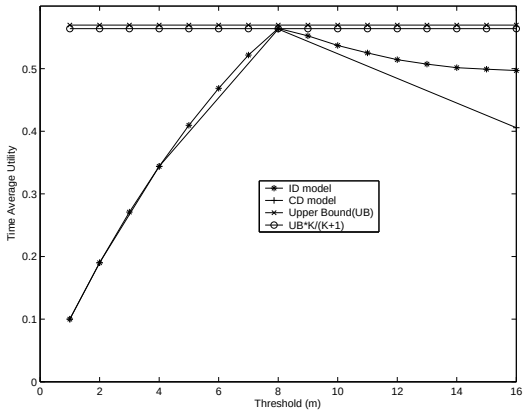

Fig. 2. Time-average utility for both models ($N = 16, K = 100, \gamma = 2$).

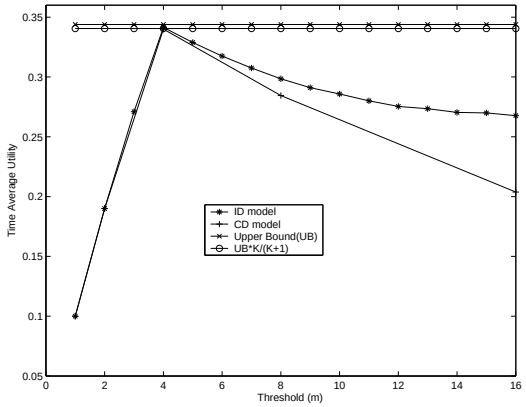

Fig. 3. Time-average utility for both models ($N = 16, K = 100, \gamma = 4$).

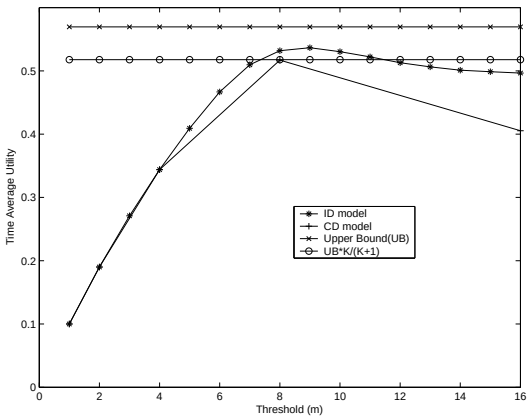

Fig. 4. Time-average utility for both models ($N = 16, K = 10, \gamma = 2$).

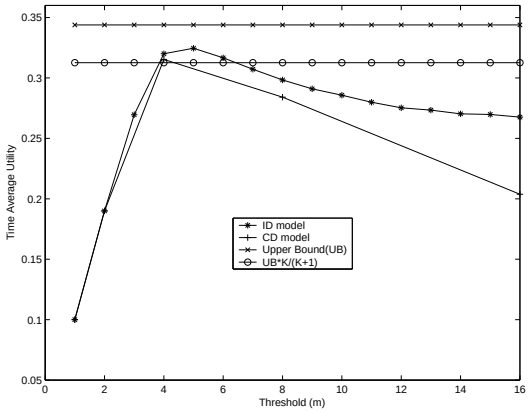

Fig. 5. Time-average utility for both models ($N = 16, K = 10, \gamma = 4$).

at the threshold of $\frac{N}{\gamma}$ is very close to that of the upper bound on the achievable performance, $U(\frac{N}{\gamma})$, particularly for large K. The time average utility at this threshold is also greater than the lower bound $\frac{K}{K+1}U(\frac{N}{\gamma})$ for both ID and CD models, as expected (Theorems 2 and 3). Also note that the performance under the ID model is better than that under the CD model, at all values of the threshold $m \in S_N$ (Theorem 4).

V. Activation Policies in a General Network Scenario

In a realistic deployment scenario, nodes may be deployed at random, and therefore, nodes will typically cover different areas in the physical space of interest. In other words, the coverage areas of two sensors may overlap only partially, or may not overlap at all. In this section, we extend our threshold activation policy to this very general scenario.

The partial coverage overlap scenario is very difficult to model and analyze, even for the special class of threshold policies. In this section, therefore, we will try to develop a solution heuristically, based on the insights obtained for the identical coverage case. We develop a threshold activation algorithm that can be implemented in a distributed manner in a general network. We then show through simulations that our solution yields a performance trend that is similar to that observed in the identical coverage case. More specifically, for an appropriately chosen threshold, our threshold activation algorithm results in a near-optimal performance even in this general network scenario.

A. Distributed Threshold Activation Policies

To motivate our distributed activation algorithm, let us assume that m_i is the target threshold for sensor i. In other words, sensor i wants to maintain a utility of $U(m_i)$ per unit area per unit time in its coverage area. In other words, if the coverage area of the sensor is denoted by $\mathcal{A}_i$, then the sensor targets to derive a utility of $|\mathcal{A}_i|U(m_i)$ per unit time. When the sensor is available (i.e., has non-zero energy), then at any decision instant, the sensor computes the current utility per unit time in its coverage area. If the current utility is less than the targeted utility, then the node activates itself; otherwise, the node remains in the available state until the next decision instant.

A sensor can compute the utility derived from its coverage area in the following manner. For a generic area element $A \in \mathcal{A}_i$, let $n(A,t)$ denote the number of sensors covering A at time t. Then the utility per unit time in the coverage area of node i is calculated as

$$\int_{\mathcal{A}_i} U(n(A,t))\, dA\ . \tag{3}$$

Assume that node i can communicate with all nodes whose coverage areas overlap with its own coverage area. Then the

sensor can periodically poll those neighbors to know their activation state. Assuming that the sensor i knows the coverage patterns of those neighbors, it can compute the targeted utility by evaluating the expression in (3). Therefore, the proposed algorithm can be realized in a distributed setting based only on local information. In practice the decision interval needs to be chosen carefully to ensure that not too much energy is wasted in the available state by periodic wakeup and polling, while guaranteeing good performance.

Note that the algorithm is motivated by the threshold activation policy discussed in the previous section, and in the case of identical sensor coverages, it reduces to a distributed implementation of the threshold policy described earlier.

B. Choice of Threshold

The thresholds m_i can be defined globally or locally, and accordingly we have two variants of our policy:

- *Global threshold policy:* In this case, the $m_i = m \ \forall i$, where the fixed threshold m is chosen appropriately.
- *Local threshold policy:* In this case, the m_i can be different for each i, depending on the local neighborhood of the individual sensor nodes.

In Section V-D, we comment on the appropriate choice of the local and global thresholds, to yield optimum performance. We can intuitively expect the local threshold policy to perform better, particularly in scenarios where there is a high spatial variance in the density of nodes in the deployment region. For the local threshold policy, nodes in areas with larger density can have a higher threshold, while nodes in a sparser region can set its threshold to a lower value. However, if the nodes are deployed more or less uniformly, then both these policies are observed to perform very well in simulations, although local threshold policy performs slightly better.

C. Discharge and Recharge Event Models

We assume that discharge (recharge) of sensors occur through discharge (recharge) events. We consider two correlation models for these discharge (recharge) events.

i) *Independent Discharge-Recharge Event Model:* Events are assumed to occur randomly in the physical space of interest, and a sensor node gets discharged (by a quantum) only when an event occurs within its coverage area. Events are assumed to occur according to a Poisson process, and are uniformly distributed in the area of interest. Thus an active sensor gets discharged of one quantum of energy if an event occurs within its coverage area, otherwise not. The recharge process is modeled similar to the discharge process. Thus "recharge events" occur randomly in the space of interest, and a sensor gets recharged by one quantum if a recharge event occurs within its coverage area.

ii) *Block-correlated Discharge-Recharge Event Model:* Here the network is divided into virtual blocks of equal sizes. Here too, discharge (recharge) events occur according to a poisson process, and are assumed to be uniformly distributed in the area of interest. However, a discharge (recharge) event occurring anywhere in the block affects all the sensors located in this block in a similar manner. This introduces spatial correlation between the discharge (recharge) times of the sensors. The degree of spatial correlation depends on the sizes of the blocks.

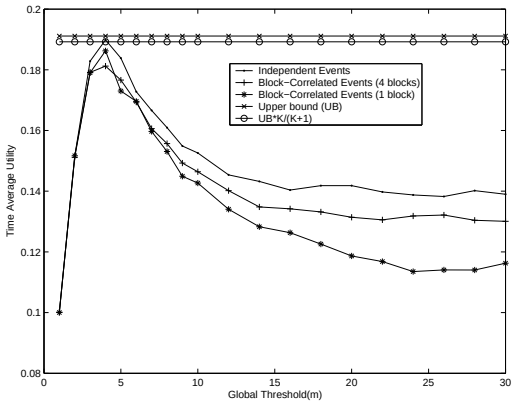

Fig. 6. Performance with global thresholds

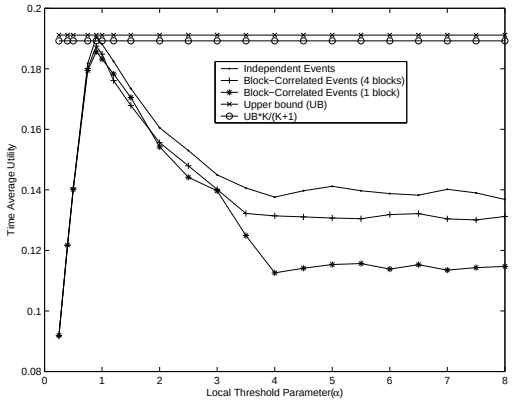

Fig. 7. Performance with local thresholds

D. Simulation Results

The performance of the node activation algorithm described above is evaluated using simulations for a wide range of parameters for both the cases of global and local thresholds. In the representative simulation results presented here, the simulation setup and the parameters used are as follows. A total of $N = 52$ sensors, each having a circular coverage pattern of radius 12 units, are thrown uniformly at random in an area of size 50×50. We use $\gamma = 2$ and $K = 100$. With these parameters, the mean coverage of the network ($\bar{N}$), defined as the average number of sensors covering any point in the deployment region, is observed to be 9.1. Therefore, $\bar{N}/\gamma \approx 2$. The utility function used is given by $U(n) = 1 - (1 - p_d)^n$, where $p_d = 0.1$.

Since the optimal policy is difficult to formulate and compute in this case, we will compare the performance of our algorithm with respect to an upper bound, computed as follows. Let A denote a generic area element in the physical space of interest, and $N(A)$ denote the number of active sensors that cover area element A. Then the following result can be proved using the same line of analysis as in the proof of Theorem 1:

Corollary 5: The optimal time-average utility for a general network of sensors is upper-bounded by

$$\int_{\mathcal{A}} U\left(\frac{N(A)}{\gamma}\right) dA\,. \tag{4}$$

Figures 6 and 7 show the performance of the various models for global and local thresholds, along with the upper bound on the performance (UB) and UB$\times\frac{K}{K+1}$. Let us define α, the local threshold parameter, as $\alpha = m_i/(\frac{n_i}{\gamma})$, where n_i is the number of sensors (including i) that cover the point where i is located. Note that in Figure 7, the time-average utility is plotted against this local threshold parameter α.

From Figure 6, we observe that the peak performance is achieved at approximately $(2\bar{N})/\gamma \approx 4)$. From Figure 6, we note that the best performance is achieved at $\alpha = 1$. A careful examination also reveals that the peak performance with local thresholds is better than the peak performance with global thresholds. Local thresholds can be computed with local information. In contrast, computation of the appropriate global threshold would require global coverage information. Thus, the use of local thresholds (over global thresholds) is justifiable from both the perspectives of overall performance and communication complexity.

Note that in all cases, the peak performance is very close to the upper bound. For the block-correlated event models, note that the degree of spatial correlation increases as the number of blocks decreases (i.e., the size of each block increases). Therefore, the figures demonstrate that the performance of threshold policies degrade as the degree of spatial correlation increases. The performance drop is particularly significant at higher values of the threshold.

REFERENCES

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam and E. Cayirci, "Wireless Sensor Networks: A Survey," *Computer Networks* 38, pp. 393-422, 2002.

[2] Bertsekas, D.P., *Dynamic Programming and Optimal Control, Volume II*, Athena Scientific, Belmont, Massachusetts, 2001.

[3] D. Blough and P. Santi, "Investigating Upper Bounds on Network Lifetime Extension for Cell-Based Energy Conservation Techniques in Stationary Ad Hoc Networks," Proc. of Mobicom 2002, Atlanta, Sept. 2002.

[4] B. Chen, K. Jamieson, H. Balakrishnan and R. Morris, "Span: An Energy Efficient Coordination Algorithm for Topology Maintenance in Ad Hoc Wireless Networks," *Proceedings of ACM Mobicom 2001*, Rome, 2001.

[5] I. Chlamtac and C. Petrioli and J. Redi, "Energy-conserving Access Protocols for Identification Networks," *IEEE/ACM Transactions on Networking*, 7(1):51-59, Feb. 1999.

[6] A. Chokalingam and M. Zorzi, "Energy Efficiency of Media Access Protocols for Mobile Data Networks," *IEEE Transactions on Communications*, 46(11):1418-21, Nov. 1998.

[7] D. Gross, C.M. Harris, *Fundamentals of Queuing Theory*, Wiley Series in Probability and Statistics, NY, 1998.

[8] K. Kar, A. Krishnamurthy and N. Jaggi, "Dynamic Node Activation in Networks of Rechargeable Sensors," To appear in P*roceedings of Infocom 2005*, Miami, March 2005.

[9] R. Kravets and P. Krishnan, "Power management techniques for mobile communications," *Proceedings of Mobicom '98*, Dallas, October 1998.

[10] S.S. Lavenberg, "The Steady State Queuing Time Distribution for the M/G/1 Finite Capacity Queue" *Management Science* 21(5):501-506, 1975.

[11] M. Papadopouli, H. Schulzrinne, "Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices," *Proceedings of Mobihoc 2001*, Long Beach, CA, Oct. 2001.

[12] M. Shaked, J.G. Shanthikumar, *Stochastic Orders and their Applications*, Academic Press, NY, 1994.

[13] S. Singh, M. Woo, C. S. Raghavendra, "Power Aware Routing In Mobile Ad-hoc Networks," *Proceedings of Mobicom '98*, Dallas, October 1998.

[14] Y. Xu, J. Heidemann, D. Estrin, "Geography-informed Energy Conservation for Ad Hoc Routing," *Proceedings of Mobicom 2001*, Roime, July 2001.

[15] W. Ye, J. Heidemann, and D. Estrin, "Medium Access Control with Coordinated, Adaptive Sleeping for Wireless Sensor Networks," *Technical Report ISI-TR-567*, USC/Information Sciences Institute, January, 2003.

[16] R.L. Wheeden, A.Zygmund, *Measure and Integral, An Introduction to Real Analysis*, Marcel Dekker, Inc., NY, pp. 694-703, 1977.

APPENDIX I: PROOF OUTLINES OF THEOREMS

Proof Outline of Lemma 1: The proof involves concavity arguments and Jensen's Inequality [16]. Let f and p be measureable functions finite a.e. on $\mathbb{R}$. Suppose, that fp and p are integrable on $\mathbb{R}$, $p \geq 0$, and $\int p > 0$. If ϕ is convex in an interval containing the range of f, then Jensen's inequality states that:

$$\phi\left(\frac{\int_R fp}{\int_R p}\right) \leq \frac{\int_R \phi(f)p}{\int_R p}.$$

Let $n(t)$ denote the number of sensors that are discharging at time t. Since $U(.)$ is concave, substituting $\phi = U(.)$, $f = n(t)$ and $p = 1$ in the above, Jensen's Inequality implies that:

$$U\left(\frac{\int_0^T n(t)dt}{T}\right) \geq \frac{\int_0^T U(n(t))dt}{T}.$$

Since, $U(.)$ is continuous, we have:

$$\lim_{T\to\infty} U\left(\frac{\int_0^T n(t)dt}{T}\right) \geq \lim_{T\to\infty} \frac{\int_0^T U(n(t))dt}{T}.$$

Therefore, it suffices to show that

$$\lim_{T\to\infty} U\left(\frac{\int_0^T n(t)dt}{T}\right) \quad \leq \quad U\left(\frac{N}{\gamma}\right).$$

Define $\psi_i(t)$ such that $\psi_i(t) = 1$ if sensor i is discharging at time t and $\psi_i(t) = 0$, otherwise. Then, continuity of $U(.)$ also implies

$$\begin{aligned}\lim_{T\to\infty} U\left(\frac{\int_0^T n(t)dt}{T}\right) &= U\left(\lim_{T\to\infty} \frac{\int_0^T n(t)dt}{T}\right) \\ &= U\left(\lim_{T\to\infty} \frac{\int_0^T \sum_{i=1}^{i=N} \psi_i(t)dt}{T}\right)\end{aligned}$$

Since $\psi_i(t)$ is positive and bounded,

$$\begin{aligned}\lim_{T\to\infty} U\left(\frac{\int_0^T n(t)dt}{T}\right) &= U\left(\lim_{T\to\infty} \sum_{i=1}^{i=N} \frac{\int_0^T \psi_i(t)dt}{T}\right) \\ &= U\left(\sum_{i=1}^{i=N} \lim_{T\to\infty} \frac{\int_0^T \psi_i(t)dt}{T}\right).\end{aligned}$$

Further, since all sensors are identical, for any i

$$U\left(\lim_{T\to\infty} \frac{\int_0^T n(t)dt}{T}\right) \quad = \quad U\left(N \lim_{T\to\infty} \frac{\int_0^T \psi_i(t)dt}{T}\right).$$

Since the recharging is modeled as Poisson processes with parameter λ, the discharging times are assumed to be exponentially

distributed with mean $1/\mu$, and each sensor is assumed to have an energy capacity equivalent to K quanta, the fraction of time a sensor is discharging corresponds to the steady state probability that the server in an $M/M/1/K$ queue is busy. Therefore, we have

$$\lim_{T\to\infty} \frac{\int_0^T \psi_i(t)dt}{T} \quad \leq \quad \frac{\gamma^K - 1}{\gamma^{K+1} - 1}.$$

Since $\gamma \geq 1$, we have $\frac{\gamma^K-1}{\gamma^{K+1}-1} \leq \frac{1}{\gamma}$ for finite K, implying

$$\lim_{T\to\infty} \frac{\int_0^T \psi_i(t)dt}{T} \quad \leq \quad \frac{1}{\gamma}.$$

Therefore we have

$$U\left(\lim_{T\to\infty} \frac{\int_0^T n(t)dt}{T}\right) \quad \leq \quad U\left(\frac{N}{\gamma}\right)$$

This implies

$$U\left(\lim_{T\to\infty} \frac{\int_0^T n(t)dt}{T}\right) \quad \leq \quad U\left(\frac{N}{\gamma}\right)$$

$$\text{i.e., } \lim_{T\to\infty} \frac{\int_0^T U(n(t))dt}{T} \quad \leq \quad U\left(\frac{N}{\gamma}\right).$$

The above theorem implies that the time-average utility under *any* policy can not be greater that $U(\frac{N}{\gamma})$. ∎

Proof Outline of Theorem 2:

The proof is based on the analysis of a queuing system equivalent to the CD model. We describe the queuing system and then prove the desired result.

Equivalent queuing system: Each sensor is modeled as a single server queue with a buffer of size K. The charging process is modeled by a Poisson process with parameter λ. We assume that arriving quanta that find a sensor fully charged are lost. In the CD model operating under a threshold policy with parameter m, exactly m of the N sensors undergo the discharge process simultaneously. The discharge process at these m sensors will be perfectly correlated, i.e. quanta of energy are consumed from each of the m sensors at exactly the same time. Therefore the N sensors will be partitioned into $c = N/m$ groups of m sensors, and the system performance can be studied by analyzing the queuing process at a representative sensor from each group. Hence, performance of the threshold policy in the CD model is analyzed using the queuing model of a polling system shown in Figure 8. The c buffers correspond to representative sensors from each of the c batches. We assume that the server polls the buffers periodically using a LUF scheduling discipline. Note that for the CD model, the LUF scheduling discipline is equivalent to a round-robin discipline. When a buffer is polled, exactly one quantum (if available) is consumed and the discharging process takes an exponentially distributed time with mean $1/\mu$. Subsequent to polling buffer j, the server polls the other $c-1$ buffers, and re-polls buffer j after at most $c-1$ exponentially distributed service times (each with mean $1/\mu$). Therefore, with respect to buffer j, the system operates like a single server queue with vacations. Let $\pi_{i,C}$ denote the long run proportion of time exactly i of the c buffers are non-empty. The long run proportion of the time the server is busy, $1-\pi_{0,C}$, determines the time average utility of the system, $V_C(m)$. That is,

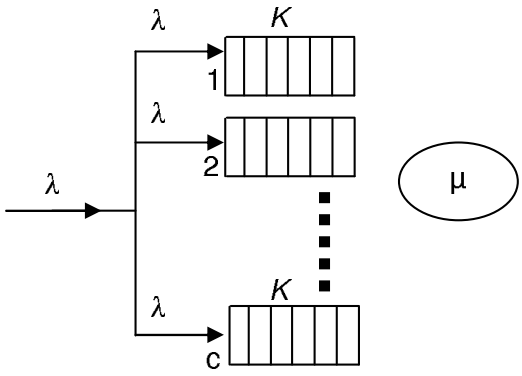

Fig. 8. Correlated Discharge Model

$$V_C(m) \quad = \quad (1-\pi_{0,C})U(m). \tag{5}$$

*Proof Outline:*We first determine an upper bound on $\pi_{0,C}$ and then show that for $m^* = N/\gamma$, $V_C(m) \geq (K/K+1)U(N/\gamma)$. To derive the bound, we analyze the single server queue with vacations as an equivalent $M/G/1/K$ queue. At this $M/G/1/K$ queue, we assume that the arrivals form Poisson process with parameter λ. We let S denote the random variable corresponding to service times at this queue. Then $S = X+Y$, where X denotes the service/discharge time (exponentially distributed with mean $1/\mu$), and Y denotes the vacation time. Note that Y is the sum of at most $c-1$ exponentially distributed random variables (each with mean $1/\mu$), depending upon whether all the other $c-1$ buffers polled were non-empty at their polling instants or not. This implies that p.d.f of Y is a mixture of Erlang distributions. We derive our desired result using a pessimistic estimate on the vacation time, i.e., we assume that all the $c-1$ buffers are always non-empty at polling instants. This implies that the p.d.f of S, $f(t)$, is Erlang distributed with c exponential stages, and that the mean service time at the $M/G/1/K$ queue is c/μ. We denote by π_0, the long run proportion of the time the server in this queue is idle. Next, we compare the performance of the $M/G/1/K$ queue with an $M/M/1/K$ queue that differs only in terms of the distribution of the service times. More precisely, if S^M denotes the service time in the $M/M/1/K$ queue, then we assume that $f^M(t) = (\mu/c)e^{-\mu t/c}$. Note that this implies that the mean service time at this $M/M/1/K$ queue is also c/μ. Further, we let π_0^M denote the long run proportion of time that the server of the $M/M/1/K$ queue is idle. The definitions of $f(t)$ and $f^M(t)$ imply that $S \leq_{st} S^M$, where $\leq_{st}$ denotes the usual stochastic ordering (Shaked and Shanthikumar [12]). Further, if $\bar{F}(t)$ and $\bar{F_M}(t)$ denote the tail probabilities of S and S^M, then $\bar{F}(t) \leq \bar{F_M}(t)$ for all $t > 0$. Then, using the analysis of Lavenberg [10] we can show that $\pi_0 \leq \pi_0^M$. Further, since the throughput of the queuing system shown in Figure 8 is at least c times the throughput of the individual $M/G/1/K$ being considered, we have $(1-\pi_{0,C})\mu \geq c(1-\pi_0)(\mu/c) \geq c(1-\pi_0^M)(\mu/c)$,

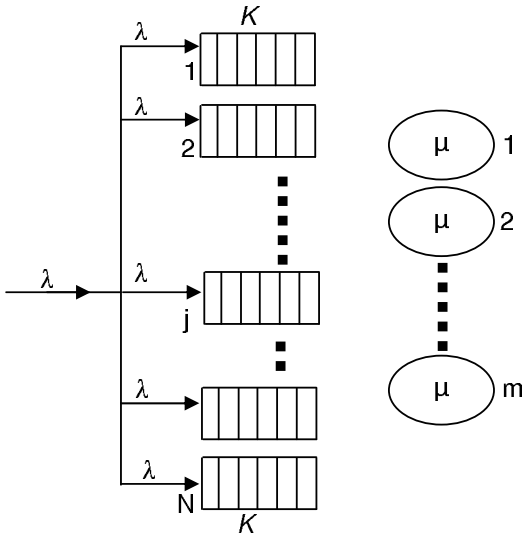

Fig. 9. Independent Discharge Model

i.e., $\pi_{0,C} \leq \pi_0^M$. From the analysis of $M/M/1/K$ queue (Gross and Harris [7]) we have $\pi_0^M = 1/\sum_{i=0}^{i=K} (c/\gamma)^i$. Now we set $m^* = N/\gamma$. Note that this implies that $c = \gamma$ and $\pi_0^M = 1/(K+1)$. Therefore,

$$\begin{aligned} V_C(m^*) &= (1-\pi_{0,C})U(N/\gamma) \\ &\geq (1-\pi_0^M)U(N/\gamma) \\ &\geq \frac{K}{K+1}U(N/\gamma) \quad \blacksquare \end{aligned} \tag{6}$$

This proves the desired result. Note that as $K \to \infty$, setting $m = N/\gamma$ yields the maximum possible time average utility under *any* policy implying that the threshold of $m = N/\gamma$ is asymptotically optimal with respect to K.

Proof Outline of Theorem 3: The proof is based on the analysis of a queuing system equivalent to the ID model. We describe the queuing system and then prove the desired result.

Equivalent queuing system: As before, each sensor is modelled as a single server queue with a buffer of size K. Under a threshold policy with parameter m in the ID model, at most m of the N sensors can be discharged simultaneously. However unlike the CD model, the discharge process at these sensors are independent. To analyze the performance of this model we construct the queuing model of a multi-server polling system (see Figure 9). Note that in this system, the buffers correspond to the N sensors. Further, the charging process is modelled by poisson arrivals and the discharging process by exponentially distributed service times. The m-server group polls the buffers periodically using a LUF scheduling policy. Therefore, with respect to a each buffer the system operates like a queuing system with vacations. When a buffer is polled, exactly one quantum (if available) is consumed and the corresponding discharging process takes an exponentially distributed time with mean $1/\mu$.

Proof Outline: The proof is similar to that used to show the result for the CD model. We let $\pi_{i,I}$ denote the long run proportion of time that exactly i of the m servers are busy. Then the time average utility of the system, $V_I(m)$ is given by

$$V_I(m) = \sum_{i=1}^{m} \pi_{i,I} U(i). \tag{7}$$

To determine a bound on $\sum_{i=1}^{m} \pi_{i,I} U(i)$ we analyze the queuing system from the perspective of a single buffer. More precisely, we analyze the queuing system with vacations using an equivalent $M/G/1/K$ queue. If $\tilde{X}$ denotes the random variable corresponding to service times of this $M/G/1/K$ queue, then $\tilde{S} = \tilde{X} + \tilde{Y}$, where $\tilde{X}$ denotes the service time – exponentially distributed with mean $1/\mu$, and $\tilde{Y}$ denotes the vacation time. (Note that $\tilde{Y}$ is a sum of at most $N-m$ exponentially distributed random variables with mean $1/m\mu$). That is, the p.d.f of $\tilde{Y}$ is a mixture of Erlang distributions. As in the analysis of the CD model, we derive our desired result using a pessimistic estimate on the vacation time, i.e., we assume that $\tilde{Y}$ has an Erlang distribution with $N-m$ stages. This implies that the p.d.f of $\tilde{S}$, $\tilde{f}(t)$ is convolution of the p.d.f of exponential and Erlang distribution. However, the mean service time at this $M/G/1/K$ queue is also c/μ ($=\frac{N-m}{m\mu} + \frac{1}{\mu} = \frac{N}{m\mu}$). We denote by $\tilde{\pi}_0$, the long run proportion of the time this server is idle. Then, we compare the performance of the $M/G/1/K$ queue with an $M/M/1/K$ queue that differs from the $M/G/1/K$ queue only in the distribution of the service times. If $\tilde{S}^M$ denotes the service time in the $M/M/1/K$ queue, then we assume that $\tilde{f}^M(t) = (\mu/c)e^{-\mu t/c}$. Further, we let $\tilde{\pi}_0^M$ denote the the long run proportion of time that the server in the $M/M/1/K$ queue is idle. Therefore we have the expectations of $\tilde{S}$ and $\tilde{S^M}$ both equal to c/μ. Further, the definitions of $\tilde{S}$ and $\tilde{S^M}$ implies that $\tilde{S} \leq_{st} \tilde{S^M}$. Then using an analysis similar to that used for the CD model, it can be shown that $\tilde{\pi}_0 \leq \tilde{\pi}_0^M$. Further, since the throughput of the queuing system shown in Figure 9 is at least N times the throughput of the individual $M/G/1/K$ being considered, we have $\sum_{i=1}^{m} i\pi_{i,I}\mu \geq N(1-\tilde{\pi}_0)(\mu/c) \geq N(1-\tilde{\pi}_0^M)(\mu/c)$. Further, the strict concavity of the utility function $U(.)$ implies that $\frac{U(i)}{i} \geq \frac{U(m)}{m}$ for all $i \leq m$. Therefore,

$$\begin{aligned} \sum_{i=1}^{m} \pi_{i,I} U(i) &\geq (U(m)/m) \sum_{i=1}^{m} i\pi_{i,I} \\ &\geq U(m)(1-\tilde{\pi}_0^M) \end{aligned}$$

From the analysis of an $M/M/1/K$ queue we can show that for $m^* = N/\gamma$, $\tilde{\pi}_0^M = 1/(K+1)$, implying that

$$V_I(m^*) \geq \frac{K}{K+1} U(N/\gamma). \quad \blacksquare \tag{8}$$

Note that as in the case of the CD model, when $K \to \infty$, setting $m = N/\gamma$ yields the maximum possible time average utility under *any* policy implying that the threshold of $m = N/\gamma$ is asymptotically optimal with respect to K.

Proof of Theorem 4:

Proof Outline: We compare the performance of the queuing system representation of the ID model under two different scheduling disciplines, LUF and g-LUF. To understand the g-LUF scheduling discipline, we focus on the system shown in

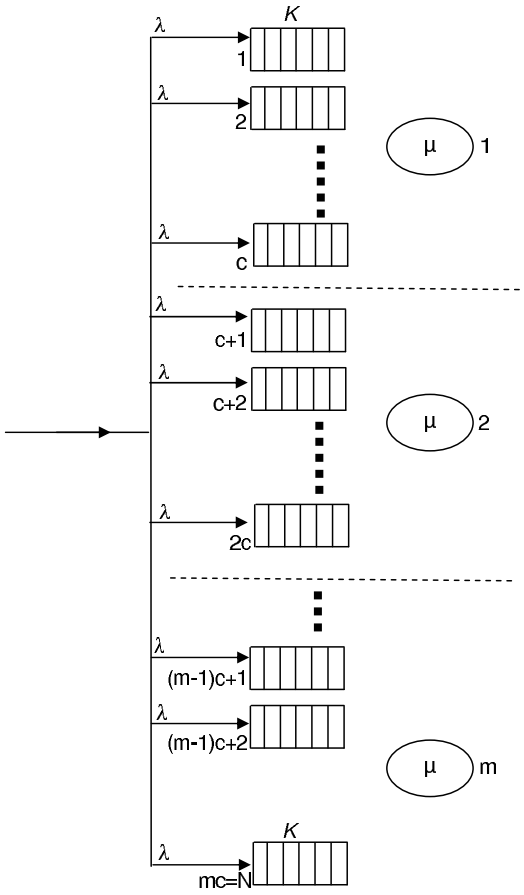

Fig. 10. Comparison of ID and CD Models

Figure 10. In this system, the N buffers are divided into m virtual groups of c buffers. A server is dedicated to each virtual group and quanta in the buffers belonging to this virtual group are served by this server exclusively. Within each virtual group, the server adopts a LUF scheduling policy. We show that when the network operates under the g-LUF scheduling discipline, the ID model yields a time average utility that is greater than that of the CD model.

We let $\pi^g_{i,I}$ denote the long run proportion of time that exactly i of the m servers are busy when the system follows the g-LUF scheduling policy. Then the time average utility of the system, $V^g_I(m)$ is given by

$$V^g_I(m) \quad = \quad \sum_{i=1}^{m} \pi^g_{i,I} U(i). \tag{9}$$

Next, we note that the stochastic behavior of the server in each virtual group is identical to that of the server in the CD model. This implies that the throughput of the system operating under the g-LUF scheduling policy, given by $\sum_{i=1}^{m} i\pi^g_{i,I}\mu = m(1-\pi_{0,C})\mu$. Further, the strict concavity of the utility function $U(.)$ implies that $\frac{U(i)}{i} \geq \frac{U(m)}{m}$ for all $i \leq m$. Therefore,

$$\begin{aligned} \sum_{i=1}^{m} \pi^g_{i,I} U(i) &\geq (U(m)/m) \sum_{i=1}^{m} i\pi^g_{i,I} \\ &\geq U(m)(1-\pi_{0,C}) \\ &\geq V_C(m). \quad \blacksquare \end{aligned} \tag{10}$$

This proves the desired result.